"Express Mail" mailing label number:

EL 830059338 US

TRANSPORT OF HIGH-BANDWIDTH DATASTREAMS OVER A NETWORK

Rajendra R. Damle
Young Lee
William C. Szeto
Robert Keith Butler
H. Michael Zadikian

CROSS-REFERENCES TO RELATED APPLICATIONS

The present invention claims priority from:

1. Provisional Patent Application Serial No. 60/268,237, entitled "EVOLUTION OF CARRRIER BACKBONE DATA NETWORK NEW REQUIREMENTS", filed February 12, 2001, and naming R. R. Damle as inventor;
2. Provisional Patent Application Serial No. 60/268,180, entitled "END-TO-END NETWORK ARCHITECTURE FOR THE NEXT GENERATION IP NETWORK", filed February 12, 2001, and naming R. K. Butler as inventor;
3. Provisional Patent Application Serial No. 60/287,973, entitled "METHOD AND APPARATUS FOR LONG-HAUL OPTICAL NETWORKING", filed April 30, 2001, and naming R. K. Butler as inventor; and
4. Provisional Patent Application Serial No. 60/295,645, entitled "TRANSPORT OF HIGH-BANDWIDTH DATASTREAMS OVER A NETWORK", filed June 4, 2001, and naming R. Damle, Y. Lee, W. Szeto, R. Butler and H. M. Zadikian as inventors.

Applicants hereby claim the benefit under 35 U.S.C. §119(e) of the foregoing-referenced provisional patent applications. The foregoing-referenced provisional patent applications are hereby incorporated by reference herein, in their entirety and for all purposes.

BACKGROUND OF THE INVENTION

Field of the Invention

This invention relates to the field of information networks, and more particularly relates to a method and system for transporting information over a network.

Description of the Related Art

Today's networks carry vast amounts of information. High bandwidth applications supported by these networks include streaming video, streaming audio, and large aggregations of voice traffic. In the future, these demands are certain to increase. To meet such demands, an increasingly popular alternative is the use of lightwave communications carried over fiber-optic cables. The use of lightwave communications provides several benefits, including high bandwidth, ease of installation, and capacity for future growth.

In an optical network, the data is transmitted between network elements (also referred to as data sources and/or sinks) over one or more communications links (or more simply, links), at least in part via an optical carrier. A general example of such a network is one in which a first network element and a second network element (e.g., optical switches or routers), both of which process and transmit data local to themselves, employ an underlying optical network to transport information from one to the other. Such an optical network provides for the transmission of data between the first network element and the second network element at least in part by use of optical carriers traversing optical fiber(s) that span the distance between network elements. A more specific example of an optical network is a network in which network components employ the SONET (Synchronous Optical NETwork) protocol.

The SONET protocol is one of a number of protocols employing an optical infrastructure. SONET is a physical transmission vehicle capable of transmission speeds in the multi-gigabit range, and is defined by a set of electrical as well as optical standards. SONET's ability to use currently-installed fiber-optic cabling, coupled with the fact that SONET significantly reduces complexity and equipment functionality

requirements, gives local and interexchange carriers incentive to employ SONET. Also attractive is the immediate savings in operational cost that this reduction in complexity provides. SONET thus allows the realization of a new generation of high-bandwidth services in a more economical manner than previously existed.

However, as the need for throughput continues to increase, even optical networks can be strained. One solution is the use of wavelength division multiplexing (WDM). In WDM systems, a number of wavelengths are used to transport the information. Further demands on the optical infrastructure can be met by increasing the data rate carried on each wavelength. This solution is not without problems.

Increasing the data rate carried on each wavelength can require newer, costly components that are designed for the higher data rates (including sophisticated modulators/receivers, newer fibers and the like). In fact, such an alternative does not necessarily provide greater total capacity per fiber (e.g., in the case of higher bit-error rates (BER)). Higher transmission rates can also limit transmission distance (due to the reduced ability to handle errors and the increased likelihood of same), and likely require more signal regeneration to overcome such problems, as well as the use of additional TDM devices. Thus, such a solution may not be well-suited to long-haul optical networking objectives (low cost, high throughput, and so on), and likely involves an increased cost per bit transmitted.

Of course, the number of wavelengths may be increased, within the capabilities of the optical technology employed. However, this can greatly complicate the management of such an optical network. For example, the complexity of an optical cross-connect increases on the order of the square of the number of entities (e.g., wavelengths) being switched. Thus, as wavelengths are added, the complexity of such a device increases dramatically.

Fig. 1 is a block diagram illustrating a network in which optical transport is employed to provide communications (e.g., data communications) between network elements 100 and 101. Network elements 100 and 101 include protocol processors 102 and 103, which provide a standard protocol stack that ensures that the information being transported is appropriately processed and routed. Protocol processors 102 and

103 operate, for example, in the electrical domain and handle high-speed datastreams 104(1)-(N) and 106(1)-(N), respectively, which are themselves in the electrical domain. Ports 108 and 110 interface electrical-optical converters 112 and 114 with protocol processors 102 and 103 by transmitting/receiving high-speed datastreams 104(1)-(N) and 106(1)-(N) to/from protocol processors 102 and 103. Those skilled in the art will appreciate that while each of electrical-optical converters 112 and 114 is depicted as having only one electrical port for simplicity, several such ports, and such ports' attendant connections with their respective protocol processors, typically exist. Those skilled in the art will also recognize that other transmission schemes can be utilized to move data between protocol processors 102 and 103, and ports 108 and 110, some of which are optical (however, for the sake of clarity such other transmission schemes are not addressed in the present discussion).

In operation, electrical-optical converters 112 and 114 convert high-speed datastreams 104(1)-(N) and 106(1)-(N) between the electrical and optical domains. The SONET protocol, in large part, defines interfaces, standards, and interactions for optical-electrical converters 112 and 114, such that each electrical-domain high-speed datastream of high-speed datastreams 104(1)-(N) and 106(1)-(N) corresponds to a datastream in the optical domain (i.e., is transmitted via an optical carrier over one of optical channels 115(1)-(N)) and vice versa.

As Fig. 1 shows, optical port groups 116 and 118 interface with optical fiber(s) 120. The term "optical port groups" is used because, in practice, each of optical channels 115(1)-(N) will have an associated optical port in both of electrical converters 112 and 114. The term "optical fiber(s)" is used due to the fact that, in practice, each of optical channels 115(1)-(N) may be transported mono-modally (i.e., via light (of a given wavelength) travelling over a mono-modal fiber), multi-modally (i.e., via light having one of a number of wavelength traveling over a multi-modal fiber), or a combination thereof (i.e., via wavelengths dispersed across some combination of mono-modal and multi-modal fibers).

It will be noted that the variable identifier "N" is used in several instances in Fig. 1 to more simply designate the final element (e.g., high-speed datastreams

104(1)-(N) and 106(1)-(N), optical channels 115(1)-(N) and so on) of a series of related or similar elements (e.g., high-speed datastreams, optical channels and so on). The repeated use of such variable identifiers is not meant to imply a correlation between the sizes of such series of elements, although such correlation may exist. The use of such variable identifiers does not require that each series of elements has the same number of elements as another series delimited by the same variable identifier. Rather, in each instance of use, the variable identified by "N" may hold the same or a different value than other instances of the same variable identifier.

Fig. 1 also illustrates that portions of electrical-optical converters 112 and 114 and optical fiber(s) 120, as well as such components' associated ports, can be viewed as an underlying network infrastructure, or (as is the case in an optical network such as that depicted in Fig. 1, among others) an installed base of fiber-optic system(s) 150. Those skilled in the art will recognize that fiber-optic system(s) 150 typically consists of many fiber-optic networks and associated subsystems which often span thousands of miles. In addition, those skilled in the art will further appreciate that fiber-optic system(s) 150 often will include miles of fiber-optic cable (e.g., which has been buried underground). Consequently, it is recognized in the art that modifying of fiber-optic system(s) 150 is, from both a technical and financial perspective, an expensive proposition which is not ordinarily undertaken unless absolutely necessary.

The transmission rate for optical channels 115(1)-(N), singly or in the aggregate, is typically limited by the available bandwidth, as in any communications system. In a network such as that depicted in Fig. 1, the chief limitation on bandwidth, and so transmission rate, is the capacity of the underlying optical communications system (e.g., fiber-optic system(s) 150), which is, in turn, limited by the physical characteristics of the optical components used in the underlying optical communications system. Such physical characteristics include the frequency of the optical carrier(s), their spacing, the transmission characteristics of the optical fiber(s), the distance between network elements and other such parameters. Typically, the bandwidth available to a datastream, and so its maximum transmission rate, is a function of the frequency (or, conversely, wavelength) of the optical carrier, and the distance that optical carrier must traverse.

Relatively recently, the data communications speeds of network components operating in the electrical domain have begun to substantially exceed the data communications speeds of network components operating in the optical domain. This situation is illustrated in Fig. 2.

Fig. 2, illustrates a network substantially analogous to that shown in Fig. 1, which includes network elements 200 and 201. Network elements 200 and 201 include protocol processors 200 and 202 and their respectively associated high-speed datastreams 204 and 206. This information is carried by one of optical channels 115(1)-(N). It will be noted that only high-speed datastreams 204 and 206 are depicted in Fig. 2. High-speed datastreams 204 and 206 correspond to only one of high-speed datastreams 104(1)-(N) and 106(1)-(N) in order to simplify the following discussions. However, in Fig. 2, the communications speed (e.g., data rate) of protocol processors 200 and 202, and so high-speed datastreams 204 and 206, exceeds the physical capabilities of the underlying network (e.g., fiber-optic system(s) 150). In other words, the electrical-domain components have transmission rates in excess of the maximum transmission rate of some or all of the optical-domain network components in fiber-optic system(s) 150. Thus, the physical characteristics of the underlying optical network limit the data rate of high-speed datastreams 204 and 206, and the extra capacity of protocol processors 200 and 202 goes unused.

As noted, it is recognized in the art that modifying fiber-optic system(s) 150 is difficult from a technical standpoint and expensive from a financial standpoint. Those skilled in the art will appreciate that the situation illustrated in Fig. 2 is seen as requiring modification of fiber-optic system(s) 150. Accordingly, most existing solutions to such a problem concentrate on partially reconstructing the underlying fiber-optic systems such that the reconstructed network can support the higher data transmission rates of protocol processors 200 and 202. Fig. 3 and 4 illustrate such solutions.

Fig. 3 illustrates one approach to the problems experienced by the network depicted in Fig. 2 by providing fiber-optic system(s) 300. In Fig. 3, as in Fig. 2, high-speed datastreams 204 and 206 operate at higher transmission rates than the

corresponding ones of high-speed datastreams 104(1)-(N) and 106(1)-(N). Because, in the given technology, information transmitted at higher transmission rates cannot be transmitted as far as information transmitted at relatively lower transmission rates, optical fiber(s) 120 are divided into optical fiber(s) 310(1)-(N+1) to allow placement of repeaters 320(1)-(N). Optical fiber(s) 310(1)-(N+1) are of lengths d_1 through d_{N+1}, respectively, each of which is shorter than length d, so that the optical carriers operating at higher transmission rates need not be transmitted as far as those in Fig. 2. Repeaters 320(1)-(N) regenerate the optical signals corresponding to high-speed datastreams 204 and 206, which are depicted in Fig. 3 as being transported over optical channels 330(1)-(N+1). Due to their shorter lengths and the regeneration that is performed by repeaters 320(1)-(N), each of optical fiber(s) 310(1)-(N+1) is capable of supporting higher-frequency (i.e., shorter wavelength) optical carriers (e.g., optical channels 330(1)-(N+1)) than optical fiber(s) 120. This enables fiber-optic system(s) 300 to support higher-speed datastreams than fiber-optic system(s) 150.

The approach depicted in Fig. 3 is an expensive alternative, from a financial standpoint, in that it involves manually restructuring large portions of fiber-optic system(s) 150 to create optical fiber(s) 310(1)-(N+1), as well as the purchase, installation, management and maintenance of repeaters 320(1)-(N). The solution depicted in Fig. 3 is also involves greater technical complexity, in that such a solution increases management and maintenance issues associated with the larger number of smaller fiber-optic segments (e.g., optical fiber(s) 310(1)-(N+1)), as well as the management and maintenance associated with repeaters 320(1)-(N).

Fig. 4 illustrates another approach, in which higher-speed components are used. To increase transmission rates within fiber-optic system(s) 150, improved optical fibers and associated subsystems are installed to replace some or all of the existing optical fibers and associated subsystems of fiber-optic system(s) 150, thus resulting in fiber-optic system(s) 400. For example, fiber-optic system(s) 400 are able to handle the higher transmission rates necessitated by high-speed datastreams 204 and 206 as a result of using improved optical fiber (depicted in Fig. 4 as optical fiber(s) 410). Such new fibers, as well as associated subsystems are designed to support communications over the required distance (e.g., d) and yet still accommodate

higher-frequency (i.e., shorter wavelength) optical carriers sufficient to support the requisite higher transmission rates. Such higher-frequency optical carriers are depicted in Fig. 4 as optical channels 420.

Unfortunately, this solution is expensive from both a technological and financial standpoint, in that such a solution often involves extensive research and development efforts (including financial support therefor) in order to create the new optical technologies – a venture which may or may not prove successful. Moreover, this solution also involves the manual reconstruction of the underlying fiber-optic network, requiring the installation of the new optical fibers and subsystems.

With respect to the approaches depicted in Figs. 3 and 4, it will be apparent to those of skill in the art that, subsequent to their implementations, a large installed base of fiber-optic systems having an upper limit of transmission limited by the fastest optical channel available will still exist, albeit with a higher upper limit than that of the original fiber-optic system(s). Accordingly, when and if the speed of the protocol processors operating in the electrical domain again exceed this new upper limit, solutions analogous to those depicted in Figs. 3 and 4 will once again require implementation.

Unfortunately, the approached depicted in Figs. 3 and 4 are financially taxing and technologically complex, as noted. In addition, those skilled in the art will recognize that, in the context of such networks, a problems analogous to that and described with regard to Fig. 2 will likely arise again. In light of this, it is apparent that a need exists in the art for a method and system which will solve those problems illustrated by the network of Fig. 2, and others, without requiring such expensive, labor-intensive and complex solutions.

What is therefore needed is a simple, inexpensive technique to transport high-bandwidth datastreams across a network. Such a technique should be applicable to optical networks, and preferably support long-haul optical networks in particular.

SUMMARY OF THE INVENTION

In one embodiment of the present invention, a method for transporting information over a network is disclosed. The method includes decomposing a datastream into a number of sub-streams, and communicating the sub-streams between a first network element and a second network element of the network by transporting each one of the sub-streams over one of a number of channels. The bandwidth of the datastream is greater than a bandwidth of any one of the channels.

In another embodiment of the present invention, a method for receiving information transported over a network is disclosed. The method includes receiving a number of sub-streams and assembling the sub-streams into a reconstructed datastream. The sub-streams are created by decomposing a datastream into the sub-streams. Each of the sub-streams is transported over the network on a corresponding one of a number of channels. A bandwidth of the datastream is greater than a bandwidth of any one of the channels.

In yet another embodiment of the present invention, an apparatus for transporting information over a network is disclosed. The apparatus includes a first sub-stream management device. The first sub-stream management device includes an input configured to receive a datastream and a number of outputs. Each of the outputs is configured to output one of a number of sub-streams. Each of the sub-streams is transported over the network on a corresponding one of a number of channels. A bandwidth of the datastream is greater than a bandwidth of any one of the channels.

In still another embodiment of the present invention, an apparatus for transporting information over a network is disclosed. The apparatus includes first sub-stream management device. The first sub-stream management device includes an output configured to output a reconstructed datastream and a number of inputs. Each of the inputs is configured to receive one of a number of sub-streams, which are created by decomposing a datastream into the sub-streams. Each of the sub-streams is transported over the network on a corresponding one of a number of channels. A bandwidth of the datastream is greater than a bandwidth of any one of the channels.

The foregoing is a summary and thus contains, by necessity, simplifications, generalizations and omissions of detail; consequently, those skilled in the art will appreciate that the summary is illustrative only and is not intended to be in any way limiting. As will also be apparent to one of skill in the art, the operations disclosed herein may be implemented in a number of ways, and such changes and modifications may be made without departing from this invention and its broader aspects. Other aspects, inventive features, and advantages of the present invention, as defined solely by the claims, will become apparent in the non-limiting detailed description set forth below.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention may be better understood, and its numerous objects, features, and advantages made apparent to those skilled in the art by referencing the accompanying drawings.

Fig. 1 is a block diagram illustrating an optical communications network.

Fig. 2 is a block diagram illustrating an optical communications network similar to that of Fig. 1.

Fig. 3 is a block diagram illustrating an optical communications network employing repeaters.

Fig. 4 is a block diagram illustrating an optical communications network employing higher-speed optical components.

Fig. 5 is a block diagram illustrating an optical communications network according to embodiments of the present invention.

Fig. 6 is a block diagram illustrating another optical communications network according to embodiments of the present invention.

Fig. 7 is a block diagram illustrating another optical communications network according to embodiments of the present invention.

Fig. 8 is a flow diagram illustrating a process according to an embodiment of the present invention.

The use of the same reference symbols in different drawings indicates similar or identical items.

DETAILED DESCRIPTION OF THE INVENTION

The following is intended to provide a detailed description of an example of the invention and should not be taken to be limiting of the invention itself. Rather, any number of variations may fall within the scope of the invention which is defined in the claims following the description.

Introduction

Embodiments of the present invention provide a method and apparatus for transporting high-bandwidth (high-speed) datastreams over a network. As noted, this is of particular concern in optical networks, where the advance of technology may not keep pace with the ever-increasing bandwidth requirements of today's (and tomorrow's) high-bandwidth applications. From a physical perspective, this challenge is particularly daunting when transporting such datastreams over long-haul optical networks.

Embodiments of the present invention address the need for increased bandwidth in current networks by decomposing each high-speed datastream into multiple sub-streams, transporting the sub-streams across the network and then re-assembling the sub-streams at the high-speed datastream's destination. Because the sub-streams' data rates are only a fraction of that of the high-speed datastream, embodiments of the present invention can provide such capabilities using existing (or available) network infrastructure (e.g., fiber-optic system(s) 150). Thus, by decomposing a high-speed datastream for transport, few, if any, changes need be made to the underlying network, avoiding potentially expensive and complex re-engineering of the existing network and allowing higher transmission rates over the given distance without the use of repeaters or other new optical infrastructure.

An Example Network Architecture

Fig. 5 is a block diagram illustrating a system substantially analogous to that shown in Fig. 2. However, the system of Fig. 5 remedies certain of the problems experienced in the system of Fig. 2 via the use of techniques according to embodiments of the present invention. As noted, each high-speed datastream is decomposed into multiple sub-streams, which are then transported across the network over separate optical channels. Once the sub-streams reach the high-speed datastream's destination, the sub-streams are re-assembled. Because the bandwidth of each sub-stream are less than that of the original high-speed datastream, the bandwidth requirements that each sub-stream places on its corresponding optical channel are reduced. In fact, if a sufficient number of sub-streams are employed, the bandwidth requirements of each sub-stream can be reduced to the point where existing optical channels are able to carry one or more such sub-streams.

Referring to Fig. 5, and as similarly depicted in Figs. 1 and 2, a network 500 is illustrated in which fiber-optic system(s) 150 are employed to allow communications between network elements 510 and 520. As in Fig. 2, network elements 510 and 520 include protocol processors 202 and 203, which provide a standard protocol stack that ensures that the information being transported is appropriately processed and routed. However, as noted in regard to Fig. 2, protocol processors 202 and 203 are capable of operating at higher data rates than protocol processors 102 and 103, and thus handle high-speed datastreams 204 and 206, which themselves have greater data rates than the corresponding ones of high-speed datastreams 104(1)-(N) and 106(1)-(N).

In contrast to other approaches, network elements 510 and 520 include sub-stream management devices 530 and 540. Sub-stream management devices 530 and 540 are capable of decomposing high-speed datastream 204 into a number of sub-streams (depicted in Fig. 5 as sub-streams 550(1)-(N)). Each of sub-streams 550(1)-(N) can then be transported across fiber-optic system(s) 150. Once transported thusly, the resulting sub-streams (which are depicted in Fig. 5 as sub-streams 560(1)-(N)) are re-assembled by sub-stream management device 540, emerging as a reconstructed

datastream, high-speed datastream 206. As before, high-speed datastream 206 is then processed by protocol processors 203.

In one embodiment, no changes to fiber-optic system(s) 150 are required or made. In such embodiments, sub-streams 550(1)-(N) correspond directly to high-speed datastreams 104(1)-(N), and so are able to traverse fiber-optic system(s) 150 in the same manner. That being the case, electrical-optical converter 112 receives sub-streams 550(1)-(N) from sub-stream management device 530 at port 106 and converts sub-streams 550(1)-(N) from the electrical domain to the optical domain, appearing at optical port group 116, which interfaces with optical fiber(s) 120, as before. The sub-streams are then transported across optical fiber(s) 120 within optical channels 115(1)-(N), to optical port group 118. Electrical-optical converter 114 then converts the sub-streams from the optical domain to the electrical domain, resulting in sub-streams 560(1)-(N) at port 110. Sub-stream management device 540 then assembles sub-streams 560(1)-(N) into high-speed datastream 206, which is then processed by protocol processors 203. It will be noted that sub-stream management devices 530 and 540 are shown separately in Fig. 5A for the sake of clarity. It will be apparent to one of skill in the art that sub-stream management devices 530 and 540 can be integrated into either of protocol processors 202 and 203, or electrical-optical converters 112 and 114, respectively, with no loss of generality.

As noted, those skilled in the art will recognize that fiber-optic system(s) 150 typically consist of many fiber-optic networks and associated subsystems which often span thousands of miles. In addition, those skilled in the art will further appreciate that fiber-optic system(s) 150 often will include miles of fiber-optic cable (e.g., which has been buried underground). Consequently, it is recognized in the art that modifying of fiber-optic system(s) 150 is, both a technically complex and financially taxing proposition which is not ordinarily undertaken. Thus, embodiments of the present invention allow the use of existing infrastructure (e.g., fiber-optic system(s) 150) to support higher data rates, avoiding the aforementioned complexities and expense.

Fig. 6 is a block diagram illustrating a network 600 which is substantially analogous to that shown in Fig. 5, but shows a more specific implementation of sub-stream management devices 530 and 540. As before, network 600 includes fiber-optic system(s) 150 that allow communications between network elements 610 and 620. In network 600, sub-stream management devices 530 and 540 are implemented as super-channel framing and management devices 630 and 640, which have access to and control over the super-channel. Such super-channel framing and management devices are discussed in detail in the Provisional Patent Application entitled, "METHOD AND APPARATUS FOR WAVELENGTH CONCATENATED CHANNEL FRAMING," as previously included by reference herein. The aggregate collection of N sub-channels supportive of a high-speed datastream such as high-speed datastream 204 or 206 is referred to herein as a "super-channel."

In this embodiment, each super-channel framing and management device splits its associated electrical-domain high-speed datastream (e.g., high-speed datastreams 204 and 206) into a series of N data channels (where N is a positive integer greater than one), where each of the N data channels has a data transmission rate that can be accommodated by existing optical channels (e.g., optical channels 115(1)-(N)) within fiber-optic system(s) 150. Preferably, each super-channel framing and management device (e.g., super-channel framing and management device 630 or 640) functions substantially transparently to protocol processors 202 and 203, as well as being transparent to fiber-optic system(s) 150. Thus, the protocol processors and underlying fiber-optic system(s) should function in substantially the same manner as in existing networks, with sub-channels able to be transported over optical channels already provided by underlying fiber-optic system(s). It is also desirable to ensure that such super-channel framing and management devices have the ability to provide varying levels of reliability of transmission for various types of transmitted data.

From a more quantifiable perspective, a high-speed datastream might have a data rate of D_{HS}. This data rate can also be related to a frequency F_{HS}, with a bandwidth $\lambda_{HS} = 1/F_{HS}$. This is the data rate of the super-channel corresponding to the high-speed datastream. In the situation posited here, the bandwidth available over any one (or at least, a given) optical channel, λ_{OC}, is less than the bandwidth of the

high-speed datastream, λ_{HS}, or $\lambda_{OC} < \lambda_{HS}$. However, when converted into a super-channel (still with a bandwidth $\lambda_{SC} = \lambda_{HS}$) having multiple sub-channels, the bandwidth of a given sub-channel, λ_{sC}, is made to be equal to or less than the bandwidth available over the corresponding optical channel, λ_{OC}, or $\lambda_{sC} \leq \lambda_{OC}$. Thus, by decomposing a high-speed datastream (and so super-channel) into sub-streams (sub-channels), the data within the high-speed datastream can be transported across fiber-optic system(s) 150. Because the sub-channels can be treated as a super-channel (i.e., the sub-channels can be treated as an entity by super-channel framing and management devices), management of the sub-channels is simplified.

As with sub-stream management devices 530 and 540, super-channel framing and management devices 630 and 640 are, in certain embodiments, able to convert high-speed datastream 204 into sub-channels 650(1)-(N) (which can also be viewed as a super-channel 655 that corresponds to high-speed datastream 204), which correspond to high-speed datastreams 104(1)-(N). These data streams are, as before, able to traverse fiber-optic system(s) 150 over optical channels 115(1)-(N). Once transported, these datastreams appear at electrical-optical converter 114 as sub-channels 660(1)-(N) (which can also be viewed as a super-channel 665 that corresponds to high-speed datastream 206). Super-channel framing and management device 640 takes in sub-channels 660(1)-(N), re-assembling the sub-channels into high-speed datastream 206. The deployment of super-channel framing and management devices 630 and 640 thus allows fiber-optic system(s) 150 to handle high-speed datastreams 204 and 206 without the reconstruction of fiber-optic system(s) 150, avoiding the expense and complexity normally associated with doing so. In a manner similar to that previous noted, and for ease of illustration, super-channel framing and management devices 630 and 640 are shown interposed between protocol processors 202 and 203, and electrical-optical converters 112 and 114, respectively. In practice, super-channel framing and management devices 630 and 640 will ordinarily be embedded within or distributed between such protocol processors and/or electrical-optical converters. In such cases, the protocol processors and/or electrical-optical converters involved may be modified to some degree, in accordance with the teachings herein.

Fig. 7 is a block diagram illustrating a system substantially analogous to that shown in Fig. 6, but with super-channel framing and management devices 630 and 640 receiving high-speed datastreams 710 and 720. High-speed datastreams 710 and 720 contain, and super-channel framing and management devices 630 and 640 have access to and control over two super-channels (depicted in Fig. 7 as super-channels 730, 735, 740 and 745) and at least one datastream/channel that remains unused by either of the super-channels (depicted in Fig. 7 as a channel 750 and a channel 755).

Fig. 7 illustrates that each of super-channel framing and management devices 630 and 640 handles respective ones of super-channels 730 and 740, each of which is substantially analogous to the super-channels described in relation to Fig. 6. Super-channel 730 (super-channel 740) can be composed of, for example, sub-channels 760(1)-(M) (sub-channels 770(1)-(M)) (where M is some integer greater than 1), where each of sub-channels 760(1)-(M) (sub-channels 770(1)-(M)) is carried by an existing optical channel of fiber-optic system(s) 150 (e.g., one of optical channels 115(1)-(N)). Further shown is that each of super-channel framing and management devices 630 and 640 also handles, respectively, at least one additional super-channel (depicted in Fig. 7 as super-channels 740 and 745). Super-channels 740 and 745 are composed of sub-channels 760(M+1)-(N) (sub-channels 770(M+1)-(N)) (where M and N are some integer greater than 1). As with sub-channels 760(1)-(M) and 770(1)-(M), each of sub-channels 760(M+1)-(N) and 770(M+1)-(N) is carried by an existing optical channel of fiber-optic system(s) 150 (e.g., one of optical channels 115(1)-(N)). Each of super-channel framing and management devices 630 and 640 also has access to at least one additional existing channel (channels 750 and 755) of fiber-optic system(s) 150 which is transported across one of optical channels 115(1)-(N) that is not currently being utilized to carry a super-channel's traffic.

Fig. 8 is a flow diagram illustrating a process according to embodiments of the present invention, in which a high-speed datastream is transported across a network having a number of channels (e.g., optical channels), none of which, taken alone, is able to provide sufficient bandwidth to support the high-speed datastream. The process begins with the reception of the high-speed datastream (step 800). The high-speed datastream (e.g., high-speed datastream 204) might be, for example, received

from a protocol processor such as protocol processor 202. The high-speed datastream is then decomposed into a number of sub-streams (step 810). This division can, for example, be performed by a sub-stream management device such as sub-stream management device 530. It will be noted, however, that the high-speed datastream can be decomposed prior to processing by protocol processor, should that be preferable in the given situation. Each of the sub-streams should have a bandwidth that is equal to or less than the corresponding one of the network's channels, over which the given sub-stream is to be transported. It will be noted that, however, one or more sub-streams may have a bandwidth greater than its corresponding channel, but be compressed to meet the bandwidth requirements of the given channel.

The high-speed datastream can be decomposed into sub-streams in any one of a number of ways. For example, a simple round-robin technique may be employed, where a portion of the high-speed datastream is periodically placed in one of a number of queues, each corresponding to one of the channels. A variation of this concept that includes framing and other mechanisms is discussed in the Provisional Patent Application entitled, "METHOD AND APPARATUS FOR WAVELENGTH CONCATENATED CHANNEL FRAMING," as previously included by reference herein.

Once decomposed into sub-streams, the high-speed datastream is transported across the network by transporting each of the (possibly compressed) sub-streams over a corresponding channel (step 820), such as one of optical channels 115(1)-(N). After being transported across the network thusly, the sub-streams are re-assembled at the high-speed datastream's destination (e.g., by a process that is the mirror-image of that used to decompose the high-speed datastream) (step 830). This can be accomplished, for example, using a sub-stream management device such as sub-stream management device 540. The high-speed datastream (e.g., high-speed datastream 206), thus re-assembled, can now be provided for further processing at the destination (step 840). This might be, for example, processing by a protocol processor such as protocol processor 203.

As noted, Fig. 8 depicts a flow diagram illustrating a process according to an embodiment of the present invention. It is appreciated that operations discussed herein may consist of directly entered commands by a computer system user or by steps executed by application specific hardware modules, but the preferred embodiment includes steps executed by software modules. The functionality of steps referred to herein may correspond to the functionality of modules or portions of modules.

The operations referred to herein may be modules or portions of modules (e.g., software, firmware or hardware modules). For example, although the described embodiment includes software modules and/or includes manually entered user commands, the various example modules may be application specific hardware modules. The software modules discussed herein may include script, batch or other executable files, or combinations and/or portions of such files. The software modules may include a computer program or subroutines thereof encoded on computer-readable media.

Additionally, those skilled in the art will recognize that the boundaries between modules are merely illustrative and alternative embodiments may merge modules or impose an alternative decomposition of functionality of modules. For example, the modules discussed herein may be decomposed into submodules to be executed as multiple computer processes, and, optionally, on multiple computers. Moreover, alternative embodiments may combine multiple instances of a particular module or submodule. Furthermore, those skilled in the art will recognize that the operations described in example embodiment are for illustration only. Operations may be combined or the functionality of the operations may be distributed in additional operations in accordance with the invention.

Alternatively, such actions may be embodied in the structure of circuitry that implements such functionality, such as the micro-code of a complex instruction set computer (CISC), firmware programmed into programmable or erasable/programmable devices, the configuration of a field-programmable gate array

(FPGA), the design of a gate array or full-custom application-specific integrated circuit (ASIC), or the like.

Each of the blocks of the flow diagram may be executed by a module (e.g., a software module) or a portion of a module or a computer system user using, for example, a computer system such as the storage router previously mentioned, or a similar network element, as well as a computer system such as computer system 210. Thus, the above described method, the operations thereof and modules therefor may be executed on a computer system configured to execute the operations of the method and/or may be executed from computer-readable media. The method may be embodied in a machine-readable and/or computer-readable medium for configuring a computer system to execute the method. Thus, the software modules may be stored within and/or transmitted to a computer system memory to configure the computer system to perform the functions of the module.

Such a computer system normally processes information according to a program (a list of internally stored instructions such as a particular application program and/or an operating system) and produces resultant output information via I/O devices. A computer process typically includes an executing (running) program or portion of a program, current program values and state information, and the resources used by the operating system to manage the execution of the process. A parent process may spawn other, child processes to help perform the overall functionality of the parent process. Because the parent process specifically spawns the child processes to perform a portion of the overall functionality of the parent process, the functions performed by child processes (and grandchild processes, etc.) may sometimes be described as being performed by the parent process.

Such a computer system typically includes multiple computer processes executing "concurrently." Often, a computer system includes a single processing unit which is capable of supporting many active processes alternately. Although multiple processes may appear to be executing concurrently, at any given point in time only one process is actually executed by the single processing unit. By rapidly changing the process executing, a computer system gives the appearance of concurrent process

execution. The ability of a computer system to multiplex the computer system's resources among multiple processes in various stages of execution is called multitasking. Systems with multiple processing units, which by definition can support true concurrent processing, are called multiprocessing systems. Active processes are often referred to as executing concurrently when such processes are executed in a multitasking and/or a multiprocessing environment.

The software modules described herein may be received by such a computer system, for example, from computer readable media. The computer readable media may be permanently, removably or remotely coupled to the computer system. The computer readable media may non-exclusively include, for example, any number of the following: magnetic storage media including disk and tape storage media. optical storage media such as compact disk media (e.g., CD-ROM, CD-R, etc.) and digital video disk storage media. nonvolatile memory storage memory including semiconductor-based memory units such as FLASH memory, EEPROM, EPROM, ROM or application specific integrated circuits. volatile storage media including registers, buffers or caches, main memory, RAM, and the like. and data transmission media including computer network, point-to-point telecommunication, and carrier wave transmission media. In a UNIX-based embodiment, the software modules may be embodied in a file which may be a device, a terminal, a local or remote file, a socket, a network connection, a signal, or other expedient of communication or state change. Other new and various types of computer-readable media may be used to store and/or transmit the software modules discussed herein.

While particular embodiments of the present invention have been shown and described, it will be obvious to those skilled in the art that, based upon the teachings herein, changes and modifications may be made without departing from this invention and its broader aspects and, therefore, the appended claims are to encompass within their scope all such changes and modifications as are within the true spirit and scope of this invention. Furthermore, it is to be understood that the invention is solely defined by the appended claims.